

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
7 March 2002 (07.03.2002)

PCT

(10) International Publication Number  
**WO 02/19077 A2**

(51) International Patent Classification?: **G06F 1/00**

(21) International Application Number: **PCT/US01/41972**

(22) International Filing Date: 31 August 2001 (31.08.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

09/653,066	1 September 2000 (01.09.2000)	US
09/711,323	9 November 2000 (09.11.2000)	US
60/278,514	23 March 2001 (23.03.2001)	US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US	09/653,066 (CIP)
Filed on	1 September 2000 (01.09.2000)
US	09/711,323 (CIP)
Filed on	9 November 2000 (09.11.2000)

(71) Applicant (for all designated States except US): **SRI INTERNATIONAL, INC. [US/US]; 333 Ravenswood Avenue, Menlo Park, CA 94025 (US)**

(72) Inventors; and

(75) Inventors/Applicants (for US only): **DE JESUS VALDES, Alfonso [US/US]; 2223 Carmelita Drive, San**

Carlos, CA 94070 (US). **SKINNER, Keith [US/US]**; 912 Suntree Court, Sunnyvale, CA 94086 (US).

(74) Agent: **FEIGENBAUM, David, L.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US)**

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— without international search report and to be republished upon receipt of that report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: PROBABILISTIC ALERT CORRELATION

(57) Abstract: This invention uses probabilistic correlation techniques to increase sensitivity, reduce false alarms, and improve alert report quality in intrusion detection systems. In one preferred embodiment, an intrusion detection system includes at least two sensors to monitor different aspects of a computer network, such as a sensor that monitors network traffic and a sensor that discovers and monitors available network resources. The sensors are correlated in that the belief state of one sensor is used to update or modify the belief state of another sensor. In another embodiment of this invention, probabilistic correlation techniques are used to organize alerts generated by different sensors in an intrusion detection system. By comparing features of each new alert with features of previous alerts, rejecting a match if a feature fails to meet or exceed a minimum similarity value, and adjusting the comparison by an expectation that certain feature values will or not match, the alerts can be grouped in an intelligent manner.

WO 02/19077 A2

## Probabilistic Alert Correlation

### **TECHNICAL FIELD**

The invention relates generally to intrusion detection, and more specifically to sensor and alert correlation in intrusion detection systems.

### **BACKGROUND**

There are three main types of intrusion detection systems currently used to detect hacker attacks on computer networks: signature analysis systems, statistical analysis systems, and systems based on probabilistic reasoning. Signature analysis systems compare current data traffic patterns with stored traffic patterns representing the signature or profile of various types of hacker attacks. These systems generate an alert if the pattern of traffic received by the network matches one of the stored attack patterns.

Statistical analysis systems compare current data traffic patterns with statistical profiles of previous traffic patterns. These systems generate an alert if a current traffic pattern is significantly different from a stored profile of "normal" traffic. Examples of both statistical- and signature-based intrusion detection systems are described in Porras, et al., "Live Traffic Analysis of TCP/IP Gateways," Internet Society's Networks and Distributed Systems Society Symposium, March 1998.

Examples of intrusion detection systems based on probabilistic reasoning are described in U.S. patent application serial number 09/653,066 entitled "Methods for Detecting and Diagnosing Abnormalities Using Real-Time Bayes Networks." In one such system, a Bayes network is established that includes models (called "hypotheses") that represent both normal traffic and attack traffic received by a computer network. Traffic actually received by the network is examined in real-time to identify its relevant characteristics or features, such as volume of data transfer, number of erroneous connection requests, nature of erroneous connection requests, ports to which connections are attempted, etc. Information about these relevant features is then provided to the Bayes network, which calculates a system belief (a probability) that the current network traffic is either normal traffic or attack traffic.

A typical intrusion detection system may include one or more sensors that monitor network traffic in the manner discussed above, and one or more other sensors that monitor network resources. A system operator or network administrator (usually a person) reviews all of the alerts generated by the system.

5 A major problem with existing intrusion detection systems is that they often provide misleading, incomplete, or low-quality information to the system operator; this may make it impossible to take proper steps to protect the network. For example, during a large-scale hacker attack each of the system's sensors may generate hundreds of alerts. Although each alert may be accurate, the sheer number of alerts could easily overwhelm the system  
10 operator. Moreover, false alarms can be triggered by normal traffic directed towards a malfunctioning network resource, such as a server. These false alarms could distract the system operator from alerts triggered by actual hacker attacks. Finally, most intrusion detection systems are unable to detect low-level attacks such as port sweeps, in which hackers slowly "probe" a network to discover its structure and weaknesses.

## SUMMARY

This invention uses probabilistic correlation techniques to increase sensitivity, reduce false alarms, and improve alert report quality in intrusion detection systems. In one preferred embodiment, an intrusion detection system includes at least two sensors to monitor different aspects of a computer network, such as a sensor that monitors network traffic and a sensor that discovers and monitors available network resources. The sensors are correlated in that the belief state of one sensor is used to update or modify the belief state of another sensor. For example, a network resource sensor may detect that a server is malfunctioning. This information is used to modify the belief state of one or more network traffic sensors so that normal network traffic directed towards the malfunctioning server does not trigger a false alarm.

In another example, a network resource sensor transmits its knowledge of network structure to a network traffic sensor. This information is used to modify the belief state of the network traffic sensor so that an attempt to communicate with a non-existent resource appears to be suspicious. By allowing different sensors to share information, the system's sensitivity to low-level attacks can be greatly increased, while at the same time greatly reducing the number of false alarms.

In another embodiment, probabilistic correlation techniques are used to organize alerts generated by different types of sensors. By comparing features of each new alert with features of previous alerts, rejecting a match if a feature fails to meet a minimum similarity value, and adjusting the comparison by an expectation that certain feature values will or will not match, the alerts can be grouped in an intelligent manner. The system operator may then be presented with a few groups of related alerts, rather than lots of individual alerts from multiple sensors.

Other features, objects, and advantages of the invention will be apparent from the description and drawings, and from the claims.

## 10 DESCRIPTION OF DRAWINGS

Figure 1 shows a preferred intrusion detection system with coupled sensors.

Figure 2 is a flowchart showing a preferred method for coupling sensors.

Figure 3 shows a preferred intrusion detection system with coupled sensors and an alert correlation device.

Figure 4 is a flowchart showing a preferred method for grouping alerts into classes of related alerts.

Figure 5 is an example of an incident class similarity matrix as used in one embodiment.

Figure 6 is an example of a display showing a correlated attack report.

Figure 7 is an example of a display showing raw sensor alerts.

Figure 8 is an example of a display showing the results of a preferred alert fusion process.

Figure 9 is an example of a display showing the results of a preferred incident matching process.

Like reference symbols in the various drawings indicate like elements.

## 25 DETAILED DESCRIPTION

### I. Introduction

In preferred embodiments, intrusion detection systems for computer networks include sensors that monitor both network traffic and network resources. Correlation techniques are

used to increase system sensitivity, reduce false alarms, and organize alerts into related groups. The sensor and alert correlation techniques described below may be used in intrusion detection systems that include any type or mix of sensors. However, probabilistic sensors similar to those described in U.S. Patent Application Serial Number 09/653,066 may 5 provide advantages over other types of sensors.

Bayes networks are preferably used to perform the probabilistic reasoning functions described below. As is more fully described in U.S. Patent Application Serial Number 09/653,066, Bayes networks include models (called "hypotheses") that represent some condition or state; probabilistic reasoning methods are used to determine a belief that a 10 current observation corresponds to one of the stored hypotheses. Before a Bayes network performs an analysis, an initial belief (called a "prior probability") regarding the various hypotheses is established. During the analysis, a new belief based on actual observations is established, and the prior probability is adjusted accordingly. In the context of the 15 embodiments described here, a relevant Bayes hypothesis might be "there is a stealthy portsweep attack against the computer network." Belief in the portsweep attack hypothesis would be strengthened if attempts to connect to several invalid or nonexistent ports are observed.

## II. Intrusion Detection with Correlated Sensors

As was discussed above, intrusion detection systems typically have several 20 independent sensors that gather different types of information. In preferred embodiments of this invention, the information gathered by the various sensors may be shared to improve overall performance of the intrusion detection system.

In a preferred intrusion detection system 100 shown in Figure 1, sensors are coupled 25 so that the belief state of one sensor (such as network resource sensor 103) affects the belief state of another sensor (such as network traffic sensor 101). Each sensor may then transmit alerts to a system operator or network administrator 105. These sensors need not be probabilistic sensors; for example, a sensor that keeps track of the number of servers in a network wouldn't need to generate a probabilistic output.

Figure 2 illustrates a preferred method by which sensors are coupled or correlated. At 30 step 201, a first (preferably probabilistic) sensor receives all or part of the belief state of a

second (not necessarily probabilistic) sensor. The belief state of the second sensor may indicate an apparent normal, degraded, or compromised state of a monitored system resource, the existence or validity of supported services, or any other relevant belief state held by a sensor in an intrusion detection system. At step 203, a prior belief state of the first sensor is 5 adjusted, the adjustment based at least in part on the second sensor's belief state. For example, the prior belief state of the first sensor or may be adjusted so that an erroneous transaction with a damaged or compromised network resource does not generate an alert. In another example, the prior belief state of the first sensor may be adjusted so that an attempted communication with a system server or resource that does not exist appears to be suspicious. 10 By allowing different sensors to share information, the system's sensitivity to low-level attacks can be greatly increased, while at the same time greatly reducing the number of false alarms.

### III. Alert Correlation

In another embodiment, probabilistic correlation techniques are used to organize 15 alerts into classes or groups of related alerts. By comparing features of each new alert with features of previous alerts, and adjusting the comparison by an expectation that certain feature values will or will not match, the alerts can be grouped in an intelligent manner. Alerts may be grouped or organized in several different ways, depending on both the type of 20 an attack and the structures of the intrusion detection system and the monitored network. For example, all alerts related to a specific attack (such as a denial of service attack) may be grouped together. Or, all alerts related to the various stages of a staged attack may be grouped together to allow the system operator to view the progression of a staged attack.

#### A. Alert Correlation Devices

Figure 3 shows a preferred intrusion detection system 300 that includes a network 25 traffic sensor 301 and network resource sensor 303. Intrusion detection system 300 may include any number of sensors, and the belief states of the sensors may be correlated in the manner discussed above. Alerts generated by sensors 301 and 303 are provided to an alert correlation device 305, which uses probabilistic reasoning (preferably Bayesian) techniques to organize alerts into classes of related alerts. These alert classes are then provided to a system 30 operator or network administrator 307.

B. Measurement of Similarity

To group alerts in an intelligent manner, it is necessary to define and determine the degree of similarity between a new alert and an existing alert or alert class. This measurement of similarity is preferably calculated by comparing similarity among shared features (also called "measures") of the alert and alert class. Examples of features to consider when assessing similarity between a new alert and an existing alert class include source IP address, destination IP address and port, type of alert, type of attack, etc. The nature of an alert may change the expectation of which features should be similar. It is therefore important to define a measurement of similarity to take into account which features are candidates for matching.

In the following discussion, X is defined to be the value of a feature or measure in a new alert, and Y is the value of that feature or measure in an existing alert class to which it is being compared. For features that have a single value, the features of a new alert and an existing alert class are defined to be similar if their feature values are equal. That is:

$$15 \quad SIM(X, Y) = \begin{cases} 1.0, & X = Y \\ 0, & \text{otherwise} \end{cases}$$

The more general case is one in which a feature of an alert includes a list of observed values, and a "hit count" of the number of times each value has been observed. For reasons of normalization, the hit count may be converted to a probability. In this sense, X and Y are lists of feature and probability values, possibly of different lengths. A probability vector describes a pattern over observed categories, where:

$p_X(C)$  = probability of category C in list X,

$p_Y(C)$  = probability of category C in list Y,

$\mathbf{P}_X$  = probability vector over categories observed for X,

$\mathbf{P}_Y$  = probability vector over categories observed for Y.

The notation  $C \in X$  is used to denote that category C occurs in list X. The similarity of the two lists is given by:

$$Sim(X, Y) = \frac{\left[ \sum_{C \in X \text{ AND } C \in Y} P_X(C) \times P_Y(C) \right]^2}{(\mathbf{P}_X \cdot \mathbf{P}_X)(\mathbf{P}_Y \cdot \mathbf{P}_Y)}$$

If the two lists are the same length, then this measure gives the square of the cosine between the two. This has an intuitive geometric property that is not shared by, say, the dot product as is commonly used in the pattern matching literature.

If the patterns  $(1, 0)$ ,  $(0, 1)$ , and  $(0.5, 0.5)$  are over the same two features, then:

$$\text{Sim}(\{1, 0\}, \{0, 1\}) = 0$$

$$\text{Sim}(\{0.5, 0.5\}, \{0, 1\}) = \text{Sim}(\{0.5, 0.5\}, \{1, 0\}) = 0.5$$

In other words, the first two patterns are orthogonal, and the third is halfway in between.

### C. Expectation of Similarity

As was discussed above, the nature of an alert may change an expectation of which features of a new alert and an existing alert class should be similar. For example, a syn flood is a type of attack in which the source IP address is typically forged. In this case, similarity in the source IP address would not be considered when assessing the overall similarity between a new alert triggered by a syn flood attack and an existing alert class.

In a preferred embodiment, the expectation of similarity is a feature-specific number between 0 and 1, with 1 indicating a strong expectation of a match, and 0 indicating that two otherwise similar alerts are likely not similar with respect to a specific feature.

An alert state is preferably represented as a probability vector over likely alert states. For each candidate alert state, each feature has a vector of the same length whose elements are the similarity expectation values (that is, numbers between 0 and 1) given the state. These expectation values may initially be assigned a value of about 0.6 to indicate a medium expectation of similarity, except where it is known that the expectation of similarity is either lower or higher. For example, in an access from a compromised host or a denial of service (DOS) attack, the attacker often spoofs (makes up) the source IP address. Therefore, the expectation values may initially be assigned a value of about 0.1, indicating a low expectation that the source IP address will match that of the attacker.

Based on the alert state, the similarity expectation may be dynamically generated as the weighted sum of the elements of an expectation table, with the weights from the (evolving) alert state distribution. This is itself an expectation in the statistical sense, conditioned on the belief over the present alert state. The similarity expectation is therefore general enough to cover the situation of an unambiguous alert or “call” from a signature-

based sensor, in which case the distribution over states is 1.0 at the state corresponding to the alert or call and 0 elsewhere. For example, algebraically for a feature or measure J:

$$E_J = \sum_{i \in \{alert\_states\}} BEL(i) E_J(i)$$

$E_J$  = Expected similarity for measure  $J$ , given present state.

$BEL(i)$  = Belief that the attack state is presently  $i$ .

$E_J(i)$  = Element  $i$  of the lookup table of expected similarity for measure  $J$  given state  $i$ .

5 Both the new alert and the alert class to which it is being compared each compute the similarity expectation for each feature. Feature similarity and similarity expectation are then combined to form a single value of alert similarity. This is done by combining feature similarities and normalizing by similarity expectation, over the set of common features. For example:

$$SIM(X, Y) = \frac{\sum_j E_j SIM(X_j, Y_j)}{\sum_j E_j}$$

10  $X$  = Candidate meta alert for matching

$Y$  = New alert

$j$  = Index over the alert features

$E_j$  = Expectation of similarity for feature  $j$

$X_j, Y_j$  = Values for feature  $j$  in alerts  $X$  and  $Y$ , respectively (may be list valued)

A similar definition can be formed by taking products rather than sums above, and then taking the geometric mean. This has some advantages, but can be overly influenced by a single extremely small value.

15 Where appropriate, similarity expectation may also cover list containment and subnet similarity. For example, an intrusion detection system may generate two alerts, one from a network sensor and one from a host sensor. To decide if the target addresses of these two alerts are similar, the expectation of similarity might be that the target address in the alert generated by the host sensor would be contained in the list of target addresses in the network sensor's alert. In another example, an intermediate expectation of similarity may be

generated if target addresses from two alerts do not match but appear to be from the same subnet.

#### D. Minimum Similarity Requirements

It may be possible to eliminate many of the similarity calculations discussed above if one or more features of an alert class are assigned a minimum similarity value. That is, if a feature of a new alert and the corresponding feature of an existing alert class do not match at some minimum similarity value, the new alert cannot be a member of that alert class. There is therefore no need to perform any additional similarity calculations comparing the new alert to that alert class. Certain features can be required to match exactly or approximately for a new alert to be considered as a candidate for an alert class.

#### E. Transition Models

When a new alert is encountered, all existing alert classes are preferably passed through their transition models to generate new prior belief states. Transition models attempt to assign a "time value" to alert confidence, as well as try to anticipate the next step in a staged attack. The "time value" is typically modeled as a multiplicative decay factor that reduces alert model confidence slowly over a period of days. The decay period varies by type of attack, and eventually decreases to some background level. This is used to downweight very old alert classes when there are multiple candidates to be compared with a new alert. Transition in time tends to decrease overall suspicion for a specific alert class, as well as spread suspicion over other hypotheses. That is, it tends to make the decision whether to group a new alert with an existing alert class less confident.

The ability to anticipate the next stage in a staged attack is also achieved by a transition in state. Internally, a transition matrix is maintained that expresses the probability of the next state in a staged attack, given the present state. Persistence (the fact that the next step may match the current) is explicitly captured. The transition gives a prior distribution over attack states that is assumed to hold immediately before the next alert is analyzed.

After the transition models generate new prior belief states for their respective alert classes, the similarity expectation for each feature in a new alert is updated. The similarity of the new alert to all existing alert classes is then computed. If none match above some

similarity threshold, the new alert is defined to be a new alert class. Otherwise, it is assumed that the new alert is a continuation of the best-matching alert class. By including a list of contributing alerts in the alert class structure, full details of the new alert are retained.

#### F. Alert Correlation Method

Figure 4 is a flowchart showing how alerts may be aggregated into classes of related alerts in a preferred embodiment. As was discussed above, both alerts and alert classes each have one or more distinguishing features that can be assigned different values. First, a set of potentially similar features shared by a new alert and one or more existing alert classes is identified (step 401). Next, an expectation of similarity between the features of the new alert and features of one or more existing alert classes is either generated or updated (step 403). Next, minimum similarity requirements for the features of the new alert and the features of one or more existing alert classes are generated or updated (step 404). After a comparison between the new alert and the existing alert class(es) is complete (step 405), the new alert is either associated with the existing alert class that it most closely matches (step 407), or new alert class is defined to include the new alert (step 409).

### IV. Alternative Embodiments and Experimental Results

The inventors plan to present various alternative embodiments and experimental results during a conference in late 2001. Relevant portions of a paper submitted for publication during that conference are shown below.

#### Introduction

In response to attacks and potential attacks against enterprise networks, administrators are increasingly deploying intrusion detection systems (IDSs). These systems monitor hosts, networks, critical files, and so forth, using a variety of signature and probabilistic techniques. See, Porras, P. and Neumann, P. "EMERALD: Event Monitoring Enabling Responses to Anomalous Live Disturbances," National Information Security Conference, 1997; and Valdes, A. and Skinner, S. "Adaptive, Model-based Monitoring for Cyber Attack Detection", Recent Advances in Intrusion Detection (RAID 2000), Toulouse,

France, October 2000. The use of such systems has given rise to another difficulty, namely, correlating a potentially large number of alerts from heterogeneous sensors. To the degree that this has been addressed in current systems, heuristic techniques have been used. In this paper, we describe a probabilistic approach to this critical problem.

5       The intrusion detection community is actively developing standards for the content of alert messages. See, Erlinger, M. and Stanniford, S. "Intrusion Detection Interchange Format," published by the Internet Engineering Task Force. Systems adhering to these evolving standards forward attack descriptions and supporting diagnostic data to an alert management interface (AMI), which may be remotely located and consolidating the reports  
10      of numerous sensors. In anticipation of these standards, we are developing systems for alert correlation and ranking.

In Valdes, A. and Skinner, S. "Blue Sensors, Sensor Correlation, and Alert Fusion," Recent Advances in Intrusion Detection (RAID 2000), Toulouse, France, October 2000, sensor coupling and an initial approach to probabilistic sensor correlation were introduced.  
15      At the time, we demonstrated the utility of the coupled sensors approach for false alarm reduction and improved detection performance. We also introduced a correlation approach based on feature-specific similarity functions. The correlation example cited at that time consisted of a single sensor examining a probe that was distributed in time. Since then, we have introduced some important innovations:

- 20      • Capability to comprehend heterogeneous sensors.
- Introduction of the probabilistic minimum match criterion, which prevents spurious matches on less important features from causing the system to correlate alerts. By specifying minimum match of unity, this capability causes the system to function as a rule-based correlation system. Explicitly tolerating slightly imperfect matches provides an important generalization of rule-based correlation methods.
- 25      • Hierarchy of sensor correlation: thread, incident, and correlated reports from multistep attack scenarios.

We have also gained considerable experimental experience in both live and simulated environments.

The remainder of this paper is organized as follows. We introduce our approach to alert fusion, using our defined alert template to report alerts from EMERALD and third-party sensors. We introduce notions of feature overlap, similarity, expectation of similarity, and minimum similarity. We show how these are used to achieve a hierarchy of alert correlation defined as inferred thread (within sensor), security incident (between sensor) and correlated attack report (between sensor and attack step). We then present preliminary results from the alerts generated by various EMERALD and third-party monitors.

### Sensor Correlation and Alert Fusion

In Valdes, A. and Skinner, S. "Blue Sensors, Sensor Correlation, and Alert Fusion", Recent Advances in Intrusion Detection (RAID 2000), Toulouse, France, October 2000, probabilistic methods for sensor correlation were introduced. Specifically, we considered two closely coupled sensors, a TCP misuse monitor and an asset availability monitor. Both are based on Bayes inference (see, Pearl, J. "Probabilistic Reasoning in Intelligent Systems", Morgan-Kaufmann (1988)), and the former sensor is aware of the state of the latter. Sensor coupling in this fashion is easily expressed in Bayes formalisms; specifically, the coupling is achieved by dynamically modifying priors in the TCP session monitor. The approach demonstrated two important capabilities:

- Failed accesses to services that appear to be invalid are viewed with greater suspicion than other failed accesses. This greatly increases sensitivity against certain stealth probe attacks.
- Certain types of failed access to services that appear "down" are tolerated, so that when a service fails (for malicious or accidental reasons) we do not generate a large number of false alarms.

Having demonstrated the utility of sensor fusion for improved sensitivity and false alarm reduction, we have further explored the general problem of sensor correlation. In particular, we examine the situation of multiple heterogeneous sensors on diverse platforms, considering realistic problems of late reports and clock drift.

Our correlation algorithm considers features as reported in a standard alert template. For each matchable feature, we have defined an appropriate similarity function that returns a

value between 0 and 1. New alerts are compared to a list of existing meta alerts. The overall similarity is the weighted average of the feature similarities, with expectation of similarity as the normalizing weights. If the minimum match criterion fails for any feature, the match is rejected regardless of the overall similarity. The new alert is correlated with the most similar 5 meta alert, assuming the similarity is above a specified minimum match threshold.

Otherwise, the new alert starts a new meta alert thread.

### Meta Alerts and the Alert Template

We have defined an enhanced alert template that enriches standards (such as the Intrusion Detection Interchange Format published by the IETF) with respect to content while 10 avoiding issues of language specifics. Our template includes the concept of alert thread, which is present in all EMERALD monitors and alert management components. Alert reports are part of the same thread if they are from the same sensor and refer to the same attack, possibly separated in time. Where provided, such as is the case with EMERALD monitors, the thread mechanism causes new reports for an existing thread to be treated as 15 updates, dramatically reducing clutter on a display interface that comprehends threads. For this reason, we identified the capability to reliably infer thread when the sensor does not provide one as a very desirable capability of alert correlation.

We include an “anomaly” field in addition to the “confidence” field used in the evolving IETF standard. We also include arrays to describe in greater detail the target(s) of 20 an attack (for example, the specific ports scanned). Finally, we include fields describing the sensor type and placement, enabling us to identify each unique instantiation of any sensor in the monitored domain.

This template is presently used by a diversity of sensors employing both signature and probabilistic techniques. This includes sensors that we have developed under the 25 EMERALD program, as well as prototype wrappers for ISS RealSecure and Checkpoint Firewall-1. Our early experiments indicate that diverse sensors will be able to fill this template with content that will be more useful to correlation engines than is currently available.

## Feature Similarity

Our probabilistic alert fusion approach preferably considers feature overlap, feature similarity, minimum similarity, and expectation of similarity. We maintain a list of “meta alerts” that are possibly composed of several alerts, potentially from heterogeneous sensors.

5 For two alerts (typically a new alert and a meta alert), we begin by identifying features they have in common (feature overlap). Such features include the source of the attack, the target (hosts and ports), the class of the attack, and time information. With each feature, we have a similarity function that returns a number between 0 and 1, with 1 corresponding to a perfect match. Similarity is a feature-specific function that considers such issues as

- 10
- How well do two lists overlap (for example, list of targeted ports)?
  - Is one observed value contained in the other (for example, is the target port of a denial-of-service (DOS) attack one of the ports that was the target of a recent probe)?
  - If two source addresses are different, are they likely to be from the same subnet?

For attack class similarity, we maintain a matrix of similarity between attack classes, with values of unity along the diagonal and off-diagonal values that heuristically express similarity between the corresponding attack classes. We prefer to consider attack classes rather than attack signatures, which are much more specific and numerous but may be erroneously or incompletely reported. For example, in our demonstration environment, we run a variant of mscan that probes certain sensitive ports, that is, it is of the attack class “portsweep.” Our host sensors have a specific signature for this attack and call it “mscan.” The Bayes sensor trades specificity for generalization capability and has no “mscan” model, but successfully detects this attack as a “portsweep”. These reports are considered similar ( $S = 1$ ) with respect to attack class.

25 Not all sensors produce all possible identifying features. For example, a host sensor provides process ID, while a network sensor does not. Features not common to both alerts are not considered for the overall similarity match.

The meta alert itself supports the threading concept, so we can visualize composing meta alerts from meta alerts.

### Situation-Specific Similarity Expectation

An important innovation we introduce is expectation of similarity. As with similarity, this is also between 0 and 1, and expresses our prior expectations that the feature should match if the two alerts are related, considering the specifics of each. For example, two probes from the same target might scan the same set of ports on different parts of our subnet (so expectation of matching target IP address is low). Also, some attacks such as SYN FLOOD spoof the source address, so we would allow a match with an earlier probe of the same target even if the source does not match (expectation of match for source IP is low).

We now give some examples of how expectation of similarity depends on the situation, that is, the features in the meta alert and the new alert.

If an alert from a sensor has a thread identifier that matches the list of sensor/thread identifiers for some meta alert, the alert is considered a match and fusion is done immediately. In other words, the individual sensor's determination that an alert is an update of or otherwise related to one of its own alerts overrides other considerations of alert similarity.

If the meta alert has received reports from host sensors on different hosts, we do not expect the target host feature to match. If at least one report from a network sensor has contributed to the meta alert and a host sensor alert is received, the expectation of similarity is that the target address of the latter is contained in the target list of the former.

In determining whether an exploit can be plausibly considered the next stage of an attack for which a probe was observed, we expect the target of the exploit (the features host and port) to be contained in the target host and port list of the meta alert.

Some sensors, particularly those that maintain a degree of state, report start and end times for an attack, while others can only timestamp a given alert. The former deal with time intervals, while the latter do not. Similarity in time comprehends overlap of the time intervals in the alerts considered for correlation, as well as the notion of precedence. We preferably do not penalize time similarity too far from unity if the time difference is plausibly due to clock drift.

Deciding whether the attacker is similar is somewhat more involved. In the case of an exact match of originating IP address, similarity is perfect. We assign high similarity if the subnet appears to match. In this way, a meta alert may potentially contain a list of

attacker addresses. At this point, we consider similarity based on containment. In addition, if an attacker compromises a host within our network, that host is added to the list of attacker hosts for the meta alert in question. Finally, for attack classes where the attacker's address is likely to be spoofed (for example, the Neptune attack), similarity expectation with respect to  
5 attacker address is assigned a low value.

### Minimum Similarity

Our correlation component implements not just expectation of similarity (which effectively acts as a weight vector on the features used for similarity matching) but may also enforce situation-specific minimum similarity. Certain features can be required to match  
10 exactly (minimum similarity for these is unity) or approximately (minimum similarity is less than unity, but strictly positive) for an alert to be considered as a candidate for fusion with another. Minimum expectation thus expresses necessary but not sufficient conditions for correlation.

The overall similarity between two alerts is zero if any overlapping feature matches at  
15 a value less than the minimum similarity for the feature (features for which no minimum similarity is specified are treated as having a minimum similarity of 0). Otherwise, overall similarity is preferably calculated as the weighted average of the similarities of the overlapping features, using the respective expectations of similarity as weights.

By appropriate settings of similarity expectation and minimum similarity, the  
20 correlation component achieves the following hierarchy of correlation. The system is composable in that we can deploy multiple instances to obtain correlation at different stages in the hierarchy. For example, we can infer threads (within sensor correlation) and then correlate threaded alerts from heterogeneous sensors into security incidents.

**Synthetic Threads:** For sensors that do not employ the thread concept, the  
25 correlation may synthesize threads by enforcing high minimum expectation similarity on the sensor itself (the thread must come from a single sensor) and the attack class, as well as source and target (IP and ports). We have wrapped the alert messages from a leading commercial sensor and observed that this facility reliably reconstructs threads.

In particular, by placing an aggregator component topologically close to an IDS, the pair is made robust against attacks that cause the IDS itself to flood, as described in the recent National Infrastructure Protection Center advisory 01-004.

5           **Security Incidents:** By suppressing minimum expectation of similarity on the sensor identifier, and relaxing expectation of similarity for this feature, we can fuse reports of the same incident from several heterogeneous sensors into a single incident report. In this case, we enforce a moderately high expectation of similarity on the attack class. This is not unity because different sensors may report a different attack class for the same attack. We construct a table of distances between attack classes that expresses which ones are acceptably 10 close. For security incident correlation, we enforce minimum expectations on the source and target of the attack. Using this technique, we have been able to fuse alert reports from commercial and EMERALD sensors into security incident reports.

15           **Correlated Attack Reports:** By relaxing the minimum expectation of similarity on the attack class, we are able to reconstruct various steps in a multistage attack. Each stage in an attack may itself be a correlated security incident as described above. In this fashion, it is possible to recognize a staged attack composed of, for example, a probe followed by an exploit to gain access to an internal machine, and then using that machine to launch an attack 20 against a more critical asset.

#### Feature Fusion

20           When the system decides to fuse two alerts, based on aggregate similarity across common features, the fused feature set is a superset of the features of the two alerts. Feature values in fused alerts are typically lists, so alert fusion involves list merging. For example, suppose a probe of certain ports on some range of the protected network matches in terms of the port list with an existing probe that originated from the same attacker subnet, but the 25 target hosts in the prior alert were to a different range of our network. The attacker address list has the new attacker address appended, and the lists of target hosts are merged. The port list matches and is thus unchanged.

30           Two important features are the sensor and thread identifiers of all the component alerts, so that the operator is always able to examine in detail the alerts that contribute to the meta alert report.

One additional feature is the priority of the meta alert, supported by our template and provided by EMERALD sensors. We are developing a component that estimates criticality based on the assets affected, the type of attack, the likelihood of attack success, and an administrative preference. The aggregator maintains the high water mark for this field. We  
5 are investigating approaches whereby the contributing threads are permitted to update their priority downward, computing meta alert priority as the maximum across thread priorities at any given time. This approach would permit downward revision of the meta alert priority.

The features presently considered in the probabilistic correlator component include sensor identification (identifier, location, name), alert thread, incident class, source and target  
10 IP lists, target TCP/UDP port lists, source user id, target user id, and time. Computations are only over features that overlap in the alert to be merged and the candidate meta alert into which it is to be merged. Incident signature is used as well, but with a low expectation of similarity as these vary widely across heterogeneous sensors.

If present, a thread identifier from the reporting sensor overrides other match criteria.  
15 A new alert that matches the sensor and thread of an existing meta alert is considered an update of the earlier alert.

The correlator first tries to infer a thread by looking for an exact match in sensor identification and incident class and signature. Note that alerts that are inferred to be from the same thread may be separated in time. The system attempts to infer threads even in  
20 incident and scenario operational modes.

Next the system checks that all overlapping features match at least at their minimum similarity value. Setting minimum expectation for some features to unity (not normally recommended) causes the system to behave like a heuristic system that requires exact matches on these features. Given that this criterion passes, we compute the overall similarity  
25 between the two alerts as follows:

$$\text{SIM}(X, Y) = \frac{\sum_j E_j \text{SIM}(X_j, Y_j)}{\sum_j E_j}$$

$X$  = Candidate meta alert for matching

$Y$  = New alert

$j$  = Index over the alert features

$E_j$  = Expectation of similarity for feature  $j$

$X_j, Y_j$  = Values for feature  $j$  in alerts  $X$  and  $Y$ , respectively (may be list valued)

Figure 5 is an example of an incident class similarity matrix. Incident class similarity is based on a notion of proximity, which at present is the result of our judgement. The proximity of class A to B reflects how reasonably an attack currently of incident class A may progress to class B. Note that this is not symmetric; we more strongly expect an exploit to follow a probe than the other way around. The incident classes shown in Figure 5 are from the EMERALD 602 message format, but any useful incident classes may be used. Note that some "default" classes such as "invalid" and "action logged" are reasonably proximal to most other classes. This is because the IETF standard does not require a common ontology, and reports from heterogeneous sensors for the same incident may not reliably represent this field. As such, we do not want to reject potential matches based on this field alone. For operational modes other than thread level aggregation, we do not recommend a high minimum similarity value for this field.

For two alerts that are extremely close in time, it is possible that the alerts may not be in time order. In this case, incident class similarity is the greater of  $\text{SIM}(X, Y)$  and  $\text{SIM}(Y, X)$ . Mathematically, the similarity computation for incident class can comprehend a discrete call (the alert is from one of the above classes) or a call that is a probability distribution over the above classes (as might result from a meta alert in which the contributing sensors do not agree on the class).

Most other features are potentially list-valued. For lists, the notion of similarity generally expresses the fraction of the smaller list that is contained in the larger. For source IP addresses, similarity also attempts to express the notion that the addresses in question may come from the same subnet.

Time similarity is preferably a step function that drops to 0.5 after one hour. A close match in time is expected only when the system operates in "incident" mode. Thread and scenario aggregation may be over time intervals of days.

## RESULTS

5

### Live Data

The following is an example of alert correlation over time, in this case correlating alerts that are components of a stealthy port sweep. The following is an example of one of the contributing alerts. In the interest of space, we do not include all the content that is in the alert and meta alert templates, but limit ourselves to the fields needed to illustrate the result.

10

```
Thread ID 69156 Class= portsweep    BEL (class) =  0.994 BEL(attack)=  
1.000  
2001-06-15 17:34:35 from 63.239.148.33 ports 1064 to 1066 duration=  
0.000  
15 dest IP aaa.bbb.30.117  
      3 dest ports: 12345{2} 27374{3} 139
```

15

This is a probe for three vulnerable ports on a single IP address in the protected network, and is detected by the system described in Valdes, A. and Skinner, S. "Adaptive, Model-based Monitoring for Cyber Attack Detection," Recent Advances in Intrusion Detection (RAID 2000), Toulouse, France, October 2000. The above is just a single step in a probe that apparently transpired over several days, and resulted in the following correlated meta alert.

25

```
Meta Alert Thread 248  
Source IPs source_IParray: 63.239.148.33 63.239.148.47  
  
Target IPs target_IParray: aaa.bbb.30.117 aaa.bbb.6.232 aaa.bbb.8.31  
30      aaa.bbb.1.166 aaa.bbb.7.118 aaa.bbb.28.83 aaa.bbb.19.121 aaa.bbb.21.130  
      aaa.bbb.6.194 aaa.bbb.1.114 aaa.bbb.16.150  
  
From 2001-06-15 17:34:35 to 2001-06-21 09:19:57  
correlated_alert_priority -1  
  
35 Ports target_TCP_portarray: 12345{4} 27374{4} 139{3}  
  
Number of threads 10 Threads :69156 71090 76696 84793 86412 87214 119525
```

124933 125331 126201  
Fused: PORT\_SCAN

We note that we have correlated events from two source addresses that were judged to  
5 be sufficiently similar. The attack is quite stealthy, consisting of a small number of attempted  
connections to single target hosts over a period of days. The list of thread identifiers permits  
the administrator to examine any of the stages in the attack. In this case, each attack stage is  
considered a portsweep; if the stages consisted of different attack classes, these would be  
listed under "Attack steps". Over the three week time period containing this attack, the IDS  
10 sensor processed over 200,000 sessions and generated 4439 alerts. The probabilistic  
correlation system produced 604 meta alerts.

### Experimentation Environment

The experimentation environment simulates an electronic commerce site, and  
provides Web and mail services over the Internet. The protected network is behind a  
15 firewall, and is instrumented with several host, network, and protocol monitors. Two  
machines are visible through the firewall, and two auxiliary machines are used for network  
monitors and alert management. The firewall blocks all services other than Web, FTP, and  
mail. We have simulated traffic that appears to come from a number of sources, and an  
attacker who executes certain well-known attacks. The background traffic includes  
20 legitimate anonymous and nonanonymous FTP use; the latter triggers false alarms in some  
signature sensors. There are also low-priority nuisance attacks, as well as low-level failed  
accesses to blocked ports, which trigger firewall log messages.

The attack begins with an MSCAN probe to the two machines visible through the  
firewall. Signature-based sensors on the two machines detect this and set the attack code for  
25 mscan in their alert report. The Bayes network sensor detects the attack as of the class  
"portsweep". The target port lists for these alerts match, as does the attacker address.  
Similarity with respect to the target address depends on the order in which the alerts arrive  
(the order at which the alerts arrive is not deterministic as the sensors are distributed). Until  
30 the network sensor's alert is received, we do not expect target addresses to match for distinct  
sensors, since a host sensor can typically report only its own address as a target. We expect  
the network sensor's target list to contain the targets seen by host sensors (tolerating

imperfect matches due to arrival order), and we expect the target address of subsequent host alerts to be contained in the meta alert's target host list. Therefore, regardless of arrival order, these alerts are acceptably similar with respect to target host.

The attacker next uses a CGI exploit to obtain the password file from the Web server.  
5 This exploit is detected by the host sensor. The next stage of the attack is a buffer overflow to gain access on the host providing mail service, detected by an EMERALD host sensor as well as by ISS.

Although telnet through the firewall is blocked, the attacker may telnet (using a password obtained from the earlier exploit) from the compromised internal host to the critical host, and he now does so. On that host, he uses another overflow attack to obtain root access.  
10 He is now free to change Web content.

Figure 6 shows the EMERALD alert console for this scenario, containing alert reports for the above attack steps as well as alerts due to false alarms and nuisance attacks. It is not very informative, except to show that more than 1000 alerts are received for this scenario,  
15 which lasts approximately 15 minutes, and the critical attack is effectively buried.

The majority of the alerts are from nonanonymous FTP write operations, which are allowed but trigger alerts from a commercial sensor.

Figure 7 shows the result of inferring synthetic threads for sensors that do not support the threading concept. For alert thread inference, we preferably set the minimum match for  
20 sensor identifier fields to unity, and require a high match for attacker and target parameters as well; we do not require a minimum similarity in the time field to allow inference of threads that occur slowly over time. Most of the alerts are threaded into sessions according to originating IP address, and the total number of alerts reported is reduced to about sixty. We have highlighted an alert, which is one of the nuisance TCP sweep attacks in the background traffic.  
25

For attack incident matching, we must look across sensors to determine if several alerts refer to the same security incident. We preferably suppress minimum similarity for  
30 sensors, set a moderately high similarity for attacker and target identifiers, and set a moderate minimum for time (to allow for clock drift and sensor response time differences). We may also require a match in attack class, making allowances for sensors that report different attack classes for a given attack sequence.

Figure 8 shows the result of incident matching for the attack scenario. Some of the critical steps in the attack reported by several sensors now result in a single alert. For example, the highlighted buffer overflow attack is reported by an EMERALD sensor as well as by ISS. These are fused into a single correlated incident report. The two mscan reports from EMERALD host sensors and the portscan report from the EMERALD Bayes TCP sensor are also fused into a single incident.

The final step in correlation is an attempt to reconstruct the attack scenario. We now relax attack class similarity to consider what attack steps might follow the present step. For example, it is reasonable to follow a probe with an exploit to a target "contained" in the probe. Also, if the meta alert includes a step indicating a likely root compromise, the host in question is added to the attacker's assets (that is, it is both a victim and now a potential attacker as well). In Figure 9, the correlated attack report correctly lists as attacker assets his originating host, an internal host compromised at an intermediate step, and the spoofed address of the SYN flood stage. In this fashion, we have successfully combined both the external and internal threads of the attack scenario.

It is worth pointing out that the EMERALD sensors in this experiment provide response directives which, if followed, would stop the attack in progress.

We have adapted and extended notions from the field of multisensor data fusion for alert correlation. The extensions are principally in the area of generalizing feature similarity functions to comprehend observables in the intrusion detection domain. The approach has the ability to fuse alerts for which the match is good but not perfect. The method considers  
5 appropriate fields in alert reports as features for a multivariate matching algorithm. Features that overlap are considered for the overall similarity calculation. For each matchable feature, we have defined an appropriate similarity function with range 0 (mismatch) to 1 (perfect match). Depending on the situation, we incorporate expectation of match values (which are used to compute a weighted average of similarity over the overlapping features), as well as a  
10 minimum match specification that unconditionally rejects a match if any feature fails to match at the minimum specified value. For each new alert, we compute similarity to existing meta alerts, and merge the new alert with the best matching meta alert, as long as the match passes a threshold value. We realize a reduction of one-half to two-thirds in alert volume in a live environment, and approach a fiftyfold reduction in alert volume in a simulated attack  
15 scenario.

**WHAT IS CLAIMED IS:**

**CLAIMS:**

- 1     1. A method for organizing alerts into alert classes, both the alerts and alert classes  
2 having a plurality of features, the method comprising the steps of:
  - 3         (a) receiving a new alert;
  - 4         (b) identifying a set of potentially similar features shared by the new alert and one  
5 or more existing alert classes;
  - 6         (c) updating a minimum similarity requirement for one or more features;
  - 7         (d) updating a similarity expectation for one or more features;
  - 8         (e) comparing the new alert with one or more alert classes, and either:
  - 9             (f1) associating the new alert with the existing alert class that the new alert most  
10 closely matches; or
  - 11             (f2) defining a new alert class that is associated with the new alert.
- 1     2. The method of claim 1 further comprising the step (a1) of passing each existing alert  
2 class through a transition model to generate a new prior belief state for each alert class.
- 1     3. A method for organizing alerts having a plurality of features, each feature having one  
2 or more values, the method comprising the steps of:
  - 3         (a) generating a group of feature records for a new alert, each feature record  
4 including a list of observed values for its corresponding feature;
  - 5         (b) identifying a set of potentially similar features shared by the new alert and one  
6 or more existing alert classes that are associated with previous alerts;
  - 7         (c) comparing the new alert to one or more alert classes;
  - 8         (d) rejecting a match if any feature for which a minimum similarity value has  
9 been set fails to meet or exceed the minimum similarity value;
  - 10         (e) adjusting the comparison by an expectation that certain feature values will or  
11 will not match, and either:
    - 12             (f1) associating the new alert with the existing alert class that the new alert most  
13 closely matches; or
    - 14             (f2) defining a new alert class that is associated with the new alert.

15        4. In an intrusion detection system that includes a plurality of sensors, each of which  
16 generates alerts when attacks or anomalous incidents are detected, a method for organizing  
17 the alerts comprising the steps of:

18              (a) receiving an alert;

19              (b) identifying a set of features that may be shared by the received alert and one  
20 or more existing alert classes;

21              (c) setting a minimum similarity value for one or more features or feature groups;  
22 comparing the new alert to one or more of the alert classes, and either:

23                  (d1) defining a new alert class that is associated with the received alert if any  
24 feature or feature group that has a minimum similarity value fails to meet or exceed its  
25 minimum similarity value; or

26                  (d2) associating the received alert with the existing alert class that the received  
27 alert most closely matches.

1        5. A method for organizing alerts into alert classes, both the alerts and alert classes  
2 having a plurality of features, the method comprising the steps of:

3              (a) receiving a new alert;

4              (b) identifying a set of potentially similar features shared by the new alert and one  
5 or more existing alert classes;

6              (c) updating a minimum similarity requirement for one or more features;

7              (d) comparing the new alert with one or more alert classes, and either:

8              (e1) associating the new alert with the existing alert class that the new alert most  
9 closely matches; or

10              (e2) defining a new alert class that is associated with the new alert.

1        6. A method for organizing alerts having a plurality of features, each feature having one  
2 or more values, the method comprising the steps of:

3              (a) generating a group of feature records for a new alert, each feature record  
4 including a list of observed values for its corresponding feature;

5              (b) identifying a set of potentially similar features shared by the new alert and one  
6 or more existing alert classes that are associated with previous alerts;

7              (c) comparing the new alert to one or more alert classes;

- 8           (d) rejecting a match if any feature for which a minimum similarity value has  
9       been set fails to meet or exceed the minimum similarity value, and either:  
10          (e1) associating the new alert with the existing alert class that the new alert most  
11       closely matches; or  
12          (e2) defining a new alert class that is associated with the new alert.

1/9

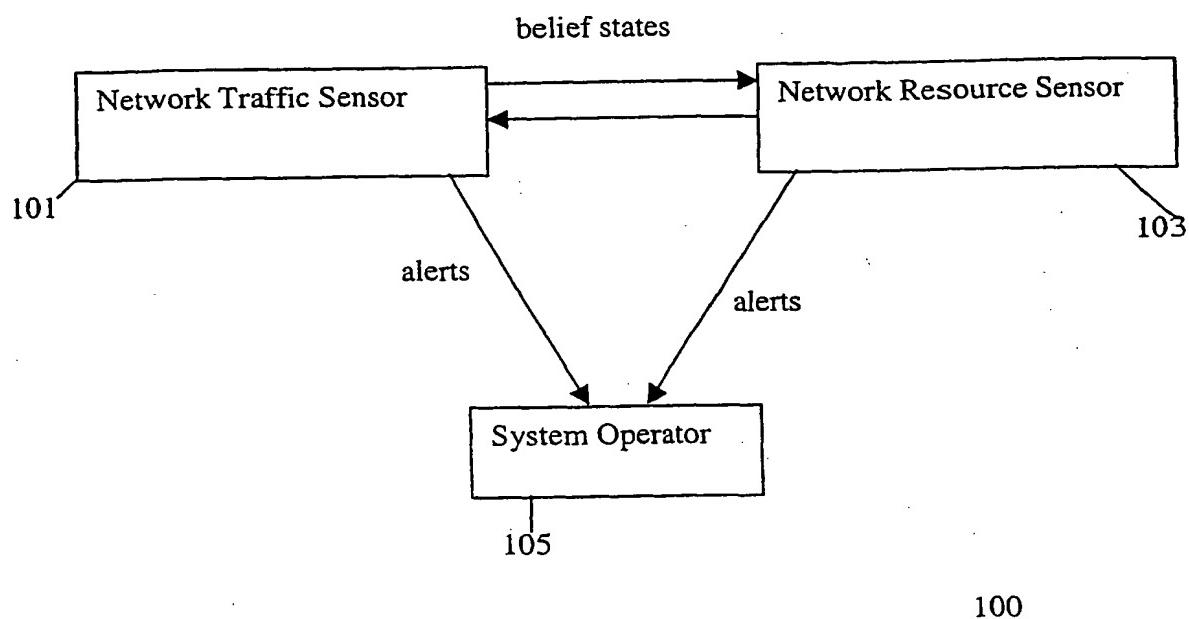
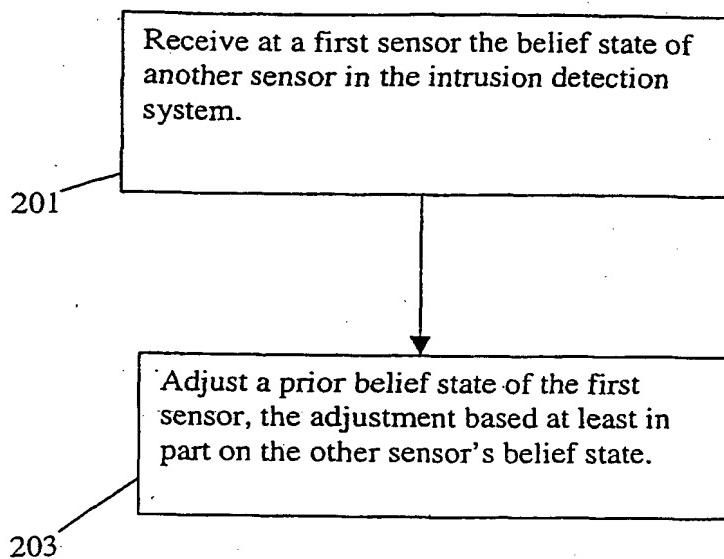


FIGURE 1



**FIGURE 2**

3/9

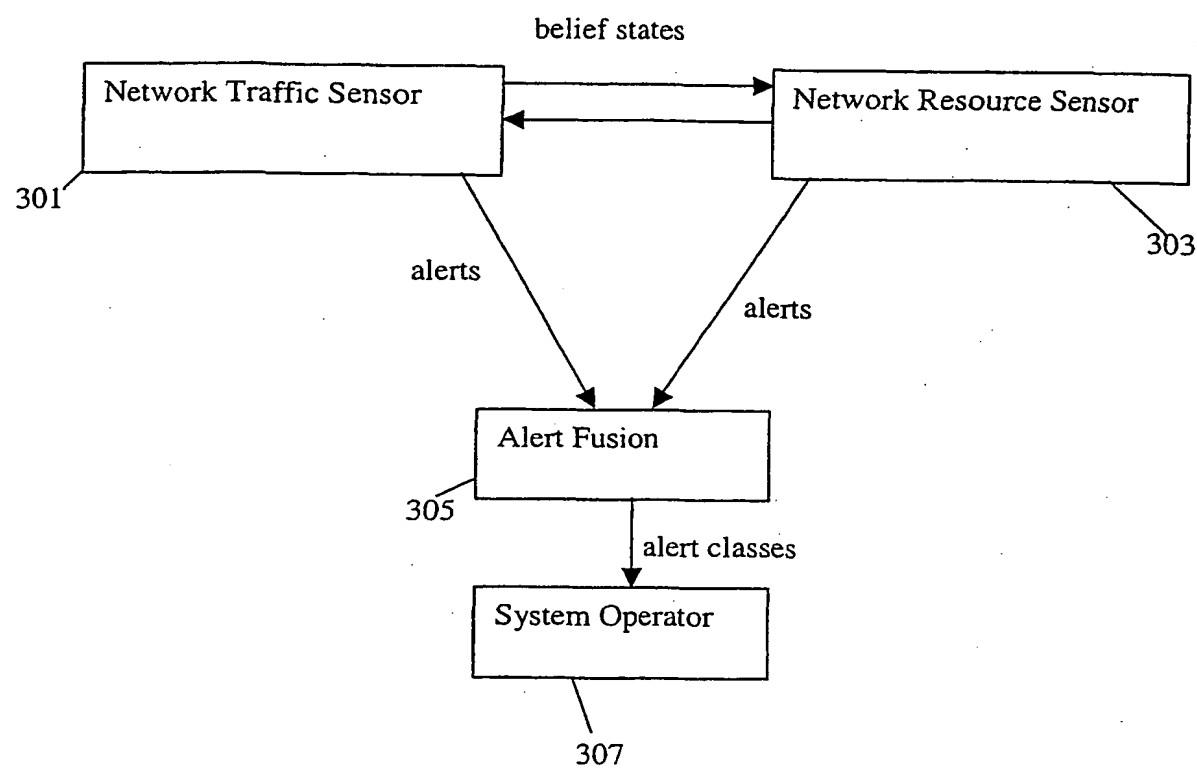


FIGURE 3

300

4/9

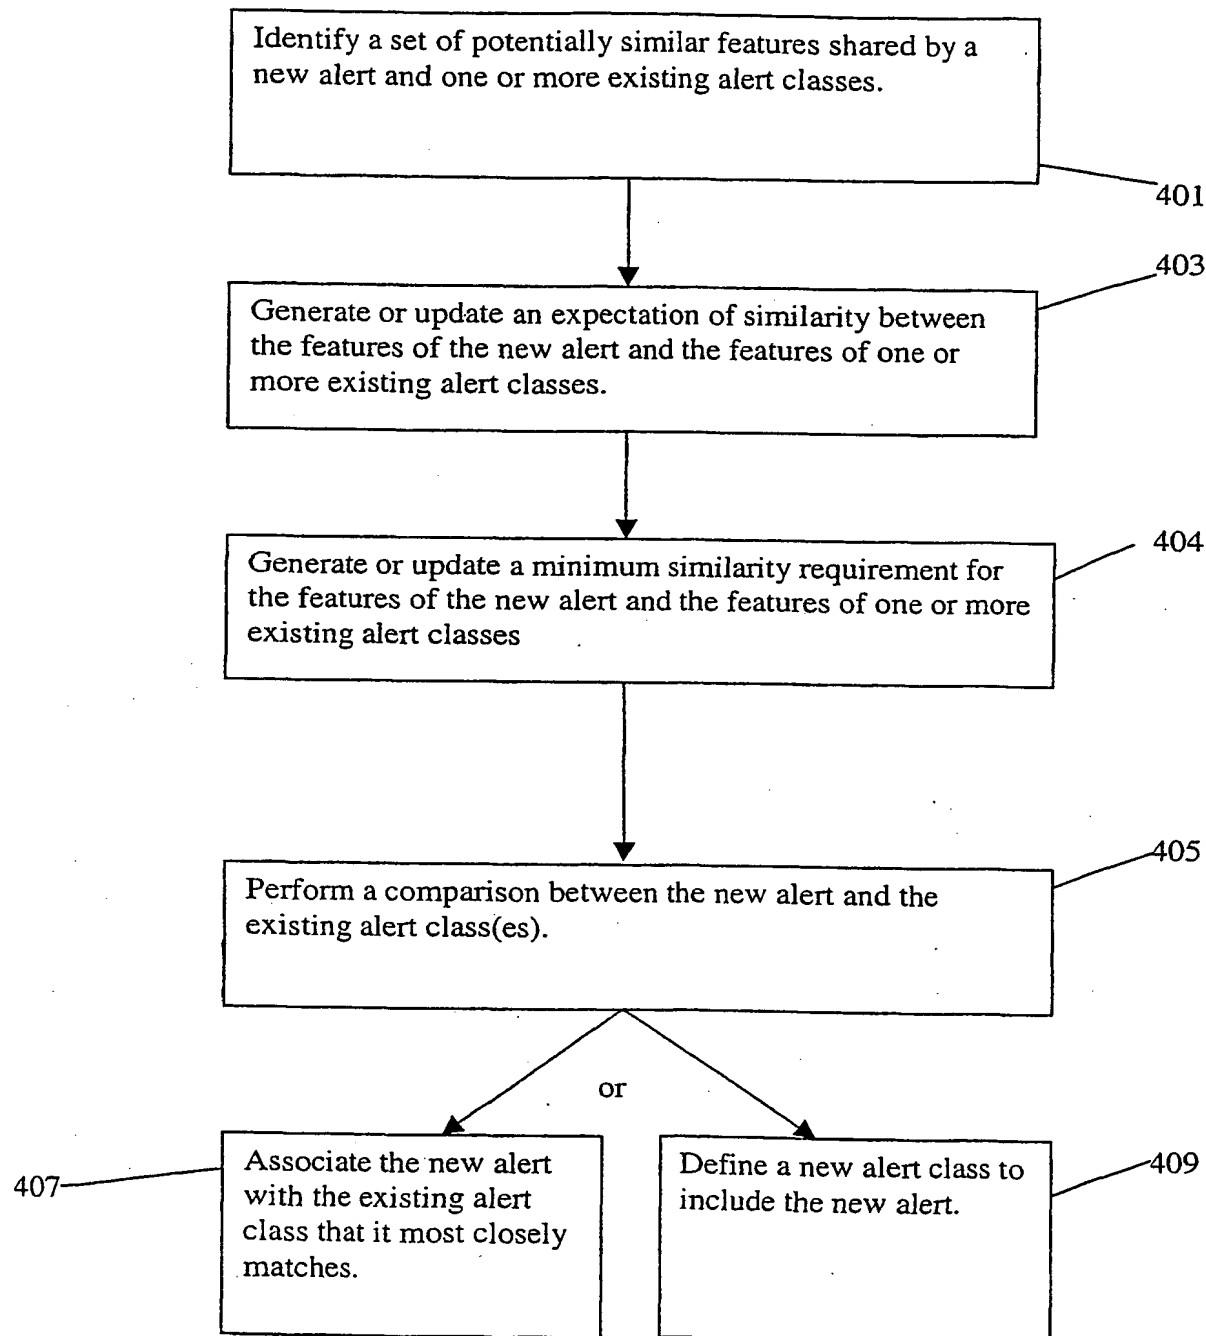


FIGURE 4

5/9

	PRIVILEGE_VIOLATION	USER_DENIAL_OF_SERVICE	INTEGRITY_VIOLATION	ACCESS_VIOLATION	SYSTEM_ENVIRONMENT_VIOLATION	USER_ENVIRONMENT_VIOLATION	ASSET_DISTRESS	SUSPICIOUS_USAGE	CONNECTION_VIOLATION	BINARY_SUBVERSION	ACTION_LOGGED
INVALID	1	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.3	0.6
PRIVILEGE_VIOLATION	0.3	1	0.6	0.3	0.6	0.6	0.6	0.4	0.3	0.4	0.1
USER_SUBVERSION	0.3	0.6	1	0.3	0.6	0.5	0.5	0.4	0.6	0.3	0.5
DENIAL_OF_SERVICE	0.3	0.3	0.3	1	0.6	0.3	0.3	0.4	0.3	0.5	0.6
PROBE	0.3	0.2	0.2	0.3	1	0.7	0.3	0.3	0.3	0.4	0.8
ACCESS_VIOLATION	0.3	0.6	0.3	0.5	0.6	1	0.6	0.3	0.3	0.4	0.1
INTEGRITY_VIOLATION	0.3	0.5	0.3	0.5	0.6	0.8	1	0.6	0.5	0.3	0.4
SYSTEM_ENVIRONMENT_VIOLATION	0.3	0.5	0.3	0.5	0.6	0.6	0.6	1	0.6	0.3	0.4
USER_ENVIRONMENT_VIOLATION	0.3	0.5	0.5	0.3	0.6	0.6	0.6	0.6	1	0.3	0.4
ASSET_DISTRESS	0.3	0.3	0.3	0.6	0.3	0.3	0.3	0.3	0.3	0.4	0.1
SUSPICIOUS_USAGE	0.3	0.3	0.5	0.3	0.5	0.6	0.5	0.6	0.5	0.3	0.6
CONNECTION_VIOLATION	0.3	0.1	0.1	0.3	0.8	0.3	0.3	0.3	0.5	0.4	1
BINARY_SUBVERSION	0.3	0.3	0.3	0.3	0.3	0.6	0.6	0.6	0.5	0.3	0.4
ACTION_LOGGED	0.3	0.3	0.3	0.3	0.6	0.5	0.3	0.3	0.3	0.4	0.3

Figure 5

6/9

**File View Tools Advanced** **Help**

**EMERALD**

Observer Name: ISS RealSecure  
 Observer Location: ntbox.emerald.sri.com  
 Observer Source: realtime  
 Local Host Time: 01/02/01 13:03:52 PST

EMERALD Development Project  
 System Design Laboratory

<b>Alert List</b>		<b>Attack Summary</b>			
Unviewed alerts 1037 Viewable alerts 1038 Hidden alerts 0 <input type="checkbox"/> Show Hidden Alerts		<b>FTP_USER:</b> FTP user command executed Date 12/08/00 15:04:43 PST End Time: 12/08/00 15:04:43 PST Class Action Logged Count 1 Updates 0 Target owl.emerald.sri.com Source 192.168.1.151 Username			
<b>Hide</b> 		<b>Other Details</b> Incident class: Action Logged signature: <b>FTP_USER</b> Alert model confidence: 70 Source TCP port 47925 Source UDP port 47925 Target TCP port 21 <small>Target UDP port 71</small>			
		<b>Recommendation</b> 			
		<b>Administrator Notes</b> 			
Acknowledgements: DARPA ITO, ISO					

Figure 6

7/9

**File View Tools Advanced** **Help**

**EMERALD**



Observer Name: eaggregate  
 Observer Location: hillsdale.cls.sri.com  
 Observer Source: realtime  
 Local Host Time: 01/02/01 13:09:13 PST



**Alert List**

Unviewed alerts 63
Viewable alerts 64
Hidden alerts 0
<input type="checkbox"/> Show Hidden Alerts

**Attack Summary** IP\_SWEEP: Fused: TCP\_ADDR\_SWEEP

Date 12/08/00 15:02:50 PST End Time: 12/08/00 15:02:50 PST
Class Probe Count 61 Updates 1
Target 130.107.12.2
Source 192.168.1.4
Username

**Other Details**

Incident class: Probe signature: TCP_ADDR_SWEEP
Alert model confidence: 100 anomaly score: 0
Target addresses: 130.107.12.2, 130.107.12.3, 130.107.12.4, 130.107.12.5, 130.107.12.6, 130.107.12.7, 130.107.12.8, 130.107.12.9, 130.107.12.10, 130.107.12.11, 130.107.12.12, 130.107.12.13, 130.107.12.14, 130.107.12.15, 130.107.12.16, 130.107.12.17, 130.107.12.18, 130.107.12.19, 130.107.12.20

**Recommendation**

Confidence level 100% that an attack was mounted from IP address 192.168.1.4
Directives:
targeted 130.107.12.2/23 130.107.12.3/23 130.107.12.4/23 130.107.12.5/23 130.107.12.6/23 130.107.12.7/23 130.107.12.8/23 130.107.12.9/23 130.107.12.10/23 130.107.12.11/23 130.107.12.12/23 130.107.12.13/23 130.107.12.14/23 130.107.12.15/23 130.107.12.16/23 130.107.12.17/23 130.107.12.18/23 130.107.12.19/23 130.107.12.20/23

**Administrator Notes**

Acknowledgements: DARPA ITO, ISO

Figure 7

8/9

File View Tools Advanced Help

# EMERALD

Observer Name: egggregate  
 Observer Location: hillsdale.cs1.sri.com  
 Observer Source: realtime  
 Local Host Time: 01/02/01 14:55:15 PST

SRI International

**EMERALD Development Project System Design Laboratory**

**Alert List**

Unviewed alerts	28
Viewable alerts	31
Hidden alerts	0
<input type="checkbox"/> Show Hidden Alerts	

**Attack Summary** BUFF\_OVER: Fused: BUFFER\_OVERFLOW->IMAP\_OVERFLOW

Date 12/08/00 14:58:51 PST End Time: 12/08/00 14:59:03 PST

Class Privilege-Violator Count 1 Updates 1

Target tigger.emerald.sri.com

Source 192.168.1.253 Username

**Other Details**

Observer IP: protocols. Outcome Generic: Unknown  
 Correlated thread ID: 418954000 observer ID: 2  
 Correlated thread ID: 0 observer ID: 0  
 Alert thread ID: 20 report ID: 329  
 Observer Type Other ID 10387 Version 1 Stream ALERT

**Recommendation**

Filter or isolate traffic stream from attacker 192.168.1.253 to victim 130.107.12.40.  
 Directives:  
 FILTER 192.168.1.253

**Administrator Notes**

Acknowledgements: DARPA ITO, ISO

Figure 8

9/9

**File View Tools Advanced Help**

**EMERALD**

**Observer Name:** eaggregate  
**Observer Location:** hillsdale.csl.sri.com  
**Observer Source:** realtime  
**Local Host Time:** 01/02/01 13:24:14 PST

**SRI International**

**EMERALD Development Project System Design Laboratory**

**Alert List**

Unviewed alerts	22
Viewable alerts	23
Hidden alerts	0
Show Hidden Alerts	<input type="checkbox"/>
3 New Alerts!	
<a href="#">Hide</a>	

BAD\_CONNECT @ 12/08 15:03   
 BAD\_CONNECT @ 12/08 15:03   
 BAD\_CONNECT @ 12/08 15:03   
 SVC\_DOWN @ 12/08 15:00   
 INTEGRITY @ 12/08 14:59   
 IP\_SWEEP @ 12/08 14:57   
 FTP\_USER @ 12/08 14:56   
 FTP\_USER @ 12/08 14:56

**Count: 3000 Updates: 78**

**Username: foo@**

**Other Details:**

Source addresses: 192.168.1.253, 130.107.12.20, 0.0.0.0 and 128.18.30.66  
 Source UDP ports: 3718, 3721, 3698, 3722 and 0  
 Source user names: foo@

**Recommendation:**

Confidence level 100% that an attack was mounted from IP address 128.18.30.66  
 Directives:  
 targeted 130.107.12.20/79 130.107.12.20/23 130.107.12.20/80 130.107.12.20/143

**Administrator Notes:**

Acknowledgements: DARPA, ITO, ISO

Figure 9

**THIS PAGE BLANK (USPTO)**

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
7 March 2002 (07.03.2002)

PCT

(10) International Publication Number  
WO 02/019077 A3

(51) International Patent Classification<sup>7</sup>: H04L 12/24

US 09/711,323 (CIP)  
Filed on 9 November 2000 (09.11.2000)

(21) International Application Number: PCT/US01/41972

(71) Applicant (for all designated States except US): SRI INTERNATIONAL, INC. [US/US]; 333 Ravenswood Avenue, Menlo Park, CA 94025 (US).

(22) International Filing Date: 31 August 2001 (31.08.2001)

(72) Inventors; and

(25) Filing Language: English

(75) Inventors/Applicants (for US only): DE JESUS VALDES, Alfonso [US/US]; 2223 Carmelita Drive, San Carlos, CA 94070 (US). SKINNER, Keith [US/US]; 912 Suntree Court, Sunnyvale, CA 94086 (US).

(26) Publication Language: English

(74) Agent: FEIGENBAUM, David, L.; Fish & Richardson P.C., 225 Franklin Street, Boston, MA 02110-2804 (US).

(30) Priority Data:

(81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC,

09/653,066 1 September 2000 (01.09.2000) US

[Continued on next page]

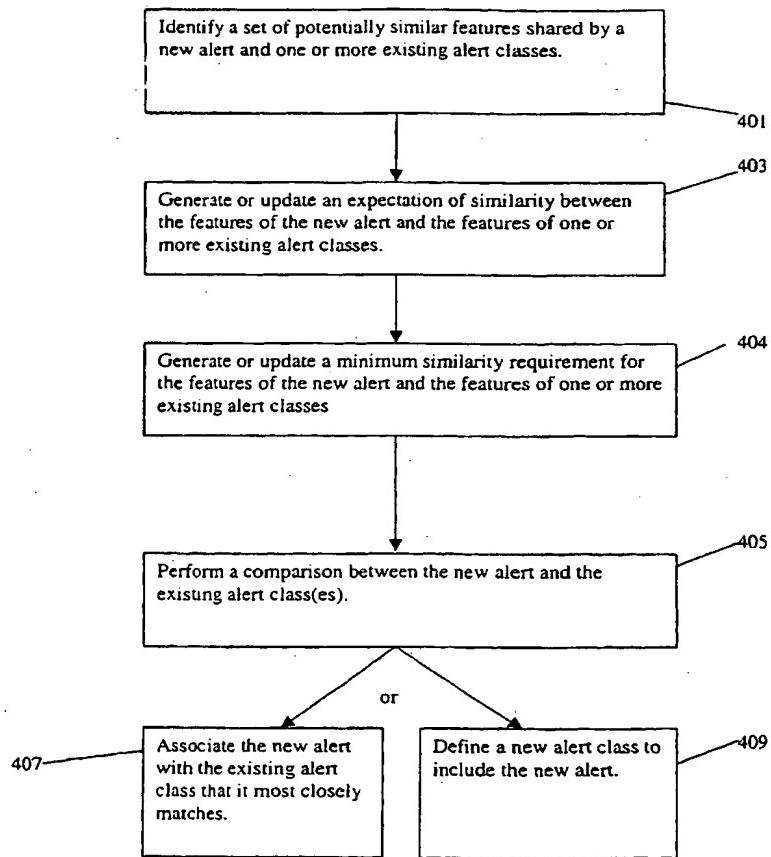
09/711,323 9 November 2000 (09.11.2000) US

60/278,514 23 March 2001 (23.03.2001) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier applications:

US 09/653,066 (CIP)  
Filed on 1 September 2000 (01.09.2000)

(54) Title: PROBABILISTIC ALERT CORRELATION



(57) Abstract: This invention uses probabilistic correlation techniques to increase sensitivity, reduce false alarms, and improve alert report quality in intrusion detection systems. In one preferred embodiment, an intrusion detection system includes at least two sensors to monitor different aspects of a computer network, such as a sensor that monitors network traffic and a sensor that discovers and monitors available network resources. The sensors are correlated in that the belief state of one sensor is used to update or modify the belief state of another sensor. In another embodiment of this invention, probabilistic correlation techniques are used to organize alerts generated by different sensors in an intrusion detection system. By comparing features of each new alert with features of previous alerts, rejecting a match if a feature fails to meet or exceed a minimum similarity value, and adjusting the comparison by an expectation that certain feature values will or not match, the alerts can be grouped in an intelligent manner.



LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, US, UZ, VN, YU, ZA, ZW.

(84) **Designated States (regional):** ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

**Published:**

— with international search report

(88) **Date of publication of the international search report:**

16 January 2003

*For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

# INTERNATIONAL SEARCH REPORT

Interr.	Application No.
PCT/US 01/41972	

**A. CLASSIFICATION OF SUBJECT MATTER**  
**IPC 7 H04L12/24**

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
**IPC 7 H04L**

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

**EPO-Internal, WPI Data, PAJ**

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 00 25527 A (ERICSSON TELEFON AB L M) 4 May 2000 (2000-05-04) abstract page 5, line 24 -page 8, line 6	1
A	WO 99 13427 A (HERRINGTON CHERYL ;ARKEL HANS VAN (US); CURTIS TERRIL J (US); GAVA) 18 March 1999 (1999-03-18) abstract page 3, line 23 -page 7, line 18 page 13, line 8 - line 18 page 17, line 17 -page 25, line 18	2-6
A	US 5 748 098 A (GRACE ANDREW) 5 May 1998 (1998-05-05) abstract column 2, line 50 -column 4, line 57	1-6

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search

**9 August 2002**

Date of mailing of the international search report

**19/08/2002**

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patenlaan 2  
NL - 2280 HV Rijswijk  
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,  
Fax: (+31-70) 340-3016

Authorized officer

**Stergiou, C**

**INTERNATIONAL SEARCH REPORT**

Information on patent family members

Internal

Application No

PCT/US 01/41972

Patent document cited in search report	Publication date		Patent family member(s)	Publication date
WO 0025527	A 04-05-2000	US AU BR WO	6253339 B1 6380799 A 9914830 A 0025527 A2	26-06-2001 15-05-2000 06-11-2001 04-05-2000
WO 9913427	A 18-03-1999	AU CA EP JP WO	9386598 A 2303107 A1 1016024 A2 2001516107 T 9913427 A2	29-03-1999 18-03-1999 05-07-2000 25-09-2001 18-03-1999
US 5748098	A 05-05-1998	AU AU CA CN DE DE DK EP ES WO HK JP JP SG	671194 B2 6111794 A 2154585 A1 1121384 A ,B 69410447 D1 69410447 T2 686336 T3 0686336 A1 2117255 T3 9419912 A1 1011162 A1 2823698 B2 8506947 T 47696 A1	15-08-1996 14-09-1994 01-09-1994 24-04-1996 25-06-1998 08-10-1998 18-01-1999 13-12-1995 01-08-1998 01-09-1994 02-07-1999 11-11-1998 23-07-1996 17-04-1998